

SYSTEM AND METHOD FOR MAPPING OF SERVICES

[1001] The present application claims priority to Provisional Application No. _____, Attorney Docket No. GRCN001/00US, entitled, "System and Method for Routing Messages Between Applications," filed March 26, 2001, which is incorporated herein by reference in its entirety.

BACKGROUND OF THE INVENTION

Field of the Invention

[1002] The present invention relates to a system and method for message routing. More specifically, the present invention relates to a message routing network for routing messages between applications.

Description of the Related Art

[1003] Corporate reliance on technology has become more complex and more pervasive. Increasingly, companies are identifying opportunities to extend their core business or cut costs using the Internet. Both trends have put increasing priority on integrating disparate business applications. For this reason, enterprise application integration (EAI) has emerged as a solution for allowing information technology departments to build bridges that are designed to unify their legacy systems into a single enterprise application. Ideally, the creation of this single enterprise application would not require sweeping changes to the underlying structures.

[1004] EAI suppliers can be viewed in four categories, by decreasing level of application independence: business process level integrators, process flow automation, data integration tools and data transport. Business process EAI offers a number of advantages over traditional middleware solutions for integrating enterprises. First, business process EAI is alleged to be application independent, allowing it to be used in any heterogeneous environment with a greater degree of reuse. Second, at its higher level of abstraction, business process EAI does not require users and implementers to have a detailed knowledge of each of the underlying technologies.

[1005] As many EAI vendors have experienced, the practice of releasing customized connectors (or adapters) for each specific enterprise software package has not proven to be scalable. Scores of adapters need to be built for each vendor (e.g., Oracle, SAP and Peoplesoft). As each supplier releases new versions of their software, EAI vendors find themselves unable to gain traction under the burden of supporting their existing adapters.

[1006] Notwithstanding the benefits of EAI, the software costs and resource investments of EAI prevent small-to-medium enterprise (SME) customers from embracing EAI solutions. For SMEs, reliance on application service providers (ASPs) represents an increasingly attractive alternative.

[1007] The ASP market is one of the fastest growing segments of the software industry. ASPs make enterprise applications (e.g., human resources administration, recruiting, travel and expense management, sales force automation) available to customers on a subscription basis. Those applications are fully managed and hosted by the ASP, providing significant cost savings to enterprises.

[1008] Some ASPs merely host and manage third-party packaged software for their customers (“managed hosters”). Others build new applications from the ground up to take advantage of the benefits and cost-savings of the Web (“webware providers”). Webware providers enjoy the profit margins and operational scalability of consumer Web companies like eBay and Yahoo, while at the same time offering the feature sets of complex enterprise software applications such as Peoplesoft and Siebel.

SUMMARY OF THE INVENTION

[1009] In accordance with the present invention, the interchange of enterprise data is supported through an open platform. This open platform can be based on a standardized interface that enables services to easily connect to and use the message interchange network. Services operating as senders, recipients, and in-transit parties can therefore leverage a framework that overlays a public network.

BRIEF DESCRIPTION OF THE DRAWINGS

[1010] The foregoing and other features and advantages of the invention will be apparent from the following, more particular description of a preferred embodiment of the invention, as illustrated in the accompanying drawings.

[1011] FIG. 1 illustrates a message exchange network.

[1012] FIG. 2 illustrates components in a message exchange network.

[1013] FIG. 3 illustrates a request response pattern.

[1014] FIG. 4 illustrates a message flow sequence.

[1015] FIG. 5 illustrates an embodiment of a message exchange network.

[1016] FIG. 6 illustrates a provisioning process.

DETAILED DESCRIPTION

[1017] An embodiment of the invention is discussed in detail below. While specific implementations are discussed, it should be understood that this is done for illustration purposes only. A person skilled in the relevant art will recognize that other components and configurations may be used without departing from the spirit and scope of the invention.

[1018] In accordance with the present invention, the interchange of enterprise data is supported through an open platform for enterprise application integration (EAI). This open platform overlays a public network (e.g., the Internet) and does not require business entities to heavily invest in specialized software and resources. As will be described in greater detail below, the present invention enables the provision of extra-enterprise application integration as a service. This service facilitates EAI efficiently and affordably to the businesses that need it the most (i.e., the small- and medium-sized enterprise (SME) market). More generally, the open platform of the present invention can be used to support services provided by business-to-business (B2B) enablers, system integrators, and other node enablers.

[1019] FIG. 1 illustrates a high-level overview of a message interchange system 100 according to the present invention. Message interchange system 100 includes a message interchange network 150 that enables SMEs 110-m, webware ASPs 120-n, and in-transit processors (ITPs) 130-p to connect to one another, integrate business processes and applications, and exchange data for mission-critical business functions. In general, ITPs 130-p are operative

to process messages that are in-transit from a sender to a recipient. ITPs 130-p can be designed to perform a variety of functions such as data transformation, enrichment, cross-reference ID mapping, filtering, credit scoring, or the like.

[1020] A directory (not shown) includes a list of all SMEs 110-m, webware ASPs 120-n and ITPs 130-p that can be accessed via message interchange network 150. Only publicly available services (i.e., those services that organizations register as accessible by any user of the network) are viewable in the directory.

[1021] In general, all applications that are connected to message interchange network 150 can be referred to as a service. In the illustrated embodiment of FIG. 1, applications owned by SMEs 110-m, webware ASPs 120-n and ITPs 130-p can each be referred to as services. Each service is owned by an organization, and an organization can have any number of services connected to message interchange network 150. The message exchange within message interchange network 150 is therefore between services.

[1022] In one embodiment, services that receive messages can take a set of arguments that further define the intended action the service will perform on a received message. For example, a service may receive the name of an operation, or may permit configuration parameters. In this environment, the service would provide a means (e.g., through a URL to documentation) for message composers to know about arguments accepted by the particular service. The message composer can then include selected arguments as a part of the service declaration in a message.

[1023] As described, services registered with message interchange network 150 represent applications that send or receive messages. An organization may, however, wish to create virtual services which act as proxies to other services. For example, a business X may have a

relationship with business Y such that messages sent to business X's service are redirected to business Y's service. Services can implement redirection through routing scripts that map invocations of the service to invocations of another service, including redirection of replies.

[1024] For each service registered by an organization with message interchange network 150, there are a number of properties and permissions that can be associated with the service. Examples include a unique service identifier, authentication information, mode of message delivery, windows of time during which messages are accepted, URL address of service, permission to invoke other services to act on a message, and rules that modify the invocation of services. These properties and permissions affect the routing of messages from or to the service.

[1025] FIG. 2 illustrates the primary functional components that operate within message interchange system 100. The five primary functional components include service software development kit (SDK) component 210, web interface component 220, message router component 230, repository component 240, and billing component 250.

[1026] SDK component 210 serves as a foundation for supported development of client applications that interface with message interchange network 150. Owning organizations can use SDK component 210 for custom integration with their software applications. As would be appreciated, SDK component 210 is not required to enable a service to access message interchange network 150. A service can use any development tool or process that would enable the service to leverage the application programming interface (API) that is supported by message router component 230.

[1027] In general, SDK component 210 enables the provision of an interface that would be available on most common platforms, and in most popular languages. In this manner, SDK

component 210 abstracts away the complex technical requirements of transferring messages using message interchange network 150.

[1028] It is a feature of the present invention that SDK component 210 need not have any business logic built into it. SDK component 210 can be used to develop plug-ins to shrink-wrapped applications, thereby greatly reducing development time. As would be appreciated, SDK component 210 can provide convenient libraries and utilities that a service may optionally use to facilitate the (1) creation and reading of messages conforming to the message router component API, and (2) authentication of users of message interchange network 150.

[1029] Repository component 240 is the primary database of message interchange network 150. Repository component 240 includes information on customer profiles, message logs, and directories. As will be described in greater detail below, message router component 230 uses repository component 240 to retrieve customer and application information that affects message routing. Message router component 230 also writes message log information to repository component 240 about messages that are processed through message interchange network 150.

[1030] Billing component 250 uses the message log information in repository component 240 to derive actual usage of message interchange network 150 by customers, and handles the invoicing and payment collection from customers. It is a feature of the present invention that the billing within message interchange system 100 can be based upon actual customer usage of message interchange network 150. For example, billing component 250 can charge customers based on a per transaction basis. In one embodiment, the per-transaction cost is based on the size of the messages being processed. As would be appreciated, these per transaction costs can be assessed to parties in a variety of ways. For example, the costs can be assessed against the

originator of the message, the intermediate services, the recipient of the message, or any combination of those parties. This billing flexibility is in sharp contrast to conventional EAI solutions that generate revenue through software license fees.

[1031] Web interface component 220 is the front-end component of message interchange network 150. Web interface component 220 interfaces directly with users by enabling login, registration, account maintenance, directory lookup, rules configuration, reporting, billing, and customer support functionality. The web interface provides online forms for data entry and can perform preliminary validations on the data. Through web interface component 220, the user can also perform queries against repository component 240 for directory lookups or reporting.

[1032] It is a feature of the present invention that message interchange network 150 is an open network architecture that not only facilitates the easy introduction of services into the network, but also enables businesses to access a robust suite of services via one connection to the message interchange network 150.

[1033] As noted, message router component 230 provides the core function of message routing and delivery within message interchange network 150. In one embodiment, message router component 230 is implemented as an Internet-based message exchange service that provides a transport level messaging service. In other words, message router component 230 need not be aware of the application semantics of a message exchange.

[1034] Thus, it is a feature of the present invention that message interchange network 150 need not inherently provide business process modeling. This is in contrast to conventional EAI solutions that may require a continual traversal up and down a protocol stack in routing a message from a sending service to a recipient service. For example, if the protocol stack

included transport, routing, transformation, and work flow layers, then each message exchange segment may require analysis and processing at each layer to determine the next service (intermediate or final) that should receive the message.

[1035] As noted, services can post messages to and retrieve messages from message router component 230 using an API. This provision of a standardized interface enables parties to easily connect to and use message interchange network 150 without being restricted in the type of message content.

[1036] In one embodiment, the protocol for posting and retrieving messages with message interchange network 150 is the Simple Object Access Protocol (SOAP). The SOAP messaging protocol defines a mechanism to pass commands and parameters between HTTP clients and servers. Through this standard object invocation protocol, HTTP is used for transport and XML is used for data encoding. The SOAP messaging protocol does not rely on the use of particular operating systems, programming languages, or object models on either the server side or the client side. As would be appreciated, other protocols can also be supported by message interchange network 150.

[1037] It is a feature of the present invention that while the message header uses extensible markup language (XML) syntax, the message body can accommodate any type of data, whether it be text or binary, encrypted or unencrypted. If the message body is also in XML form, then the message body can opt to use a schema based on an industry standard such as ebXML, BizTalk, RosettaNet, OAGIS, or any other suitable standard.

[1038] In one embodiment, message exchange through message interchange network 150 is asynchronous. Recipient services can be configured to poll message interchange network 150

for incoming messages, or, if they have their own server, can have message interchange network 150 push messages to them.

[1039] After a sending service posts a message to message interchange network 150, one or more in-transit services 130-p can operate on the message before it reaches the recipient service. In-transit services can perform useful operations on messages, such as data transformation, enrichment, cross-reference ID mapping, filtering, credit scoring, or the like. Through the standardized interface, in-transit services 130-p can independently join the message interchange network 150 and operate on messages. This flexibility encourages independent third parties to build services that can be plugged into message interchange network 150. It is a feature of the present invention that such an open network would encourage third parties to market a data service that generates revenue based upon the level of utilization of the service.

[1040] As noted, in-transit services can be included in a message path that begins at a sending service and terminates at a recipient service. As will be described in greater detail below, sending services can explicitly specify a set of services to operate on a given message. In addition, recipient services can specify services that should operate on messages before delivery to the recipient service. In one example, a recipient may always want messages to pass through a filtering service to screen out messages from unknown senders.

[1041] Messaging through message interchange network 150 can be as secure as the participants desire. Each service registered with message interchange network 150 can specify a security policy declaring encryption and authentication levels for message interchange network 150 to enforce. For messages that flow through in-transit services, a sender can also specify the permissions for each in-transit service to access or operate on parts of the message.

[1042] In one embodiment, message interchange network 150 uses the secure HTTPS protocol to support secure transport connections when a service posts a message or polls for messages, and when message interchange network 150 pushes messages to a client server. Authentication can be based on either username/password or certificates.

[1043] SSL encryption as part of HTTP can be used to provide data protection during message transmission over the public Internet. In general, this level of protection is sufficient for most situations. Services can, however, perform their own extra encryption of message documents to keep them private even from message interchange network 150. Services that add extra encryption should ensure, however, that all services that operate on the message documents have the necessary keys to decrypt the documents.

[1044] As is well known, the authentication protocol of SSL includes a server's presentation of a certificate to clients. Accordingly, message interchange network 150 presents a server certificate to services that connect for posting or polling. The connecting service has the option of then providing either a username/password or certificate for message interchange network 150 to authenticate the service. The form of client authentication is a configuration option within the profile message interchange network 150 maintains for each service.

[1045] When message interchange network 150 pushes messages to a service, the service's server should present a server certificate to message interchange network 150 for authentication of the service. For the reverse authentication, the service can then require either a username/password or certificate from message interchange network 150. Again, that option can be configured in the profile information message interchange network 150 maintains for the service.

[1046] As a message flows through a selection of services on the way to the recipient service, and as the recipient service's response returns to the sending service, message interchange network 150 maintains an audit trail of all operations on the message and all services that touched the message. The audit trail serves several purposes. First, it enables message interchange network 150 to reconstruct the message history in the case of queries on the message trail. Second, it allows message interchange network 150 to compile a usage report for any service for reporting and billing purposes.

[1047] Having described the general framework of message interchange network 150, a more detailed description of a message transaction lifecycle within message interchange network 150 is provided with reference to FIG. 3.

[1048] In this framework, a message can be embodied as a self-contained collection of information to serve a particular purpose, such as a request, a response, a notification, or an acknowledgement. As noted, message interchange network 150 can generally be agnostic about the content of a message other than header information that affects routing of the message.

[1049] In one embodiment, request, response, and notification messages can be defined. A request message expects a subsequent response message from the recipient(s) to be returned to the sender. Request messages may represent inquiries, but might also represent update requests that only expect a return status response. If an error occurs in routing a request message, message interchange network 150 returns an error response message to the sender.

[1050] A response message is issued by a recipient of a request message. The response message references the original request message. Failure of the response message may result in an error response message being returned to the sender of the original request message.

[1051] A notification message is a one-way message. No response to the notification message is expected back to the sender. Message interchange network 150 can regard any response message referencing a notification message as an invalid message. If a notification message fails, no error message is returned to the sender.

[1052] As would be appreciated, further messages can be defined for message interchange network 150. For example, a cancel message can also be defined, wherein the cancel message is used by the sender to cancel a previous message.

[1053] The operation of these messages is now described with reference to the request/response illustration of FIG. 3. This illustration demonstrates a typical example of a sending service 310, such as an enterprise, making an inquiry to a recipient service 360, such as a webware provider. In one embodiment, a sender's application 312 that connects to message interchange network 150 is a desktop application. In another embodiment, a sender's application 312 that connects to message interchange network 150 is an enterprise server, or an EAI package.

[1054] The first step in the message transaction process is the creation of a message. In one embodiment, a sender formats the messages to conform to an XML schema for messages. This XML schema prescribes the format for message headers while allowing any kind of data to be included in the message body (or payload). As part of message construction, sending service 310 specifies the recipient service(s) 360 of the message. In one embodiment, a recipient service's name includes an organization and a specific service provided by that organization. The service name can be generally represented in the message via a globally unique ID.

[1055] The actual set of elements contained in a message depend on whether the message is being posted or delivered. In one embodiment, a message includes a header element, a body element and/or attachments. In one embodiment, the attachments are based on multi-part Multipurpose Internet Mail Extensions (MIME).

[1056] An embodiment of a message that includes header and body elements is included in Appendix A. In this example, the cardinality of elements is indicated as '?' or {0:1} for an optional instance; '*' or {0:N} for zero or more instances; and '+' or {1:N} for one or more instances. No symbol or {1} represents a required single instance. As would be appreciated, the actual message format can differ depending on the protocol. In particular, protocols other than the SOAP protocol can be used.

[1057] The header element includes routing and invocation information. The header as posted by a sending service is often modified by message interchange network 150 for delivery to the receiving service.

[1058] The body element includes the documents the sender is sending to the recipient(s). These documents can also be operated upon by one or more services. As noted, the documents can be in the form of XML or any other representation, including text, binary, etc. In one embodiment, all or part of the documents that are being sent are included in an attachment to the message.

[1059] While messages will typically have a similar overall structure, the actual composition of elements can differ between the various message types and between messages as posted and as delivered. For example, some elements in a sent message can be changed or not be included

in the message as delivered, such as elements particular to constructing a route. Some elements can also be inserted only in the message as delivered, such as identifier elements.

[1060] If the sending service wishes to have the message routed through any services before delivery to the recipient service(s), the sending service can specify an explicit sequence of services that should operate on the message. The sender can also implicitly include services in the route for a message through the specification of routing scripts associated with the defining service. Routing scripts are described in greater detail below.

[1061] After a message is constructed, the message is posted to message interchange network 150. This process is illustrated in FIG. 3 as the posting of a message by application 312 to message post interface 324. As noted, in one embodiment, the posting of a message is performed using the SOAP messaging protocol.

[1062] If sending service 310 posts a message that does not have well-formed XML, the message posting is rejected and an error response is returned. In general, messages can be rejected for a variety of other reasons. For example, a message can be rejected if the service indicated in the message header as the sender is not the same as the actual sender of the message, the message is a duplicate posting of a previous message, a service attempts to reply to a message for which it was not a recipient, or a response message does not reference a prior message.

[1063] In one embodiment, each message posted by a service can have a unique handle assigned by the service to identify the message. This unique handle can be used to provide a means for message interchange network 150 to detect duplicate postings of the same message. Duplicate postings can occur in the case of failure recovery by the service. In one embodiment,

if a service desires that message interchange network 150 should reject duplicate postings of a message, then the service could provide unique handles for messages and set a “potential duplicate” flag in messages that may be a duplicate posting. It should be noted that regardless of whether or not a service provides a unique handle for a message, message interchange network 150 can assign a globally unique session identifier to each posted message.

[1064] After a message is posted, message interchange network 150 routes the message to the recipient service(s) 360. The routing of the message is based upon a route calculation by message interchange network 150. The calculated route includes all intermediary services 350 that are scheduled to operate on the message en route to recipient service(s) 360. The calculated route can be based on routing instructions specified explicitly in the message header and/or on routing scripts pre-defined by the sending service 310, recipient service 360, or any in-transit services 350 that have been included within the calculated route.

[1065] In general, routing scripts define a procedure for enabling determination of at least part of a route. This procedure can be based on any type of criteria. For example, a procedure can be defined that determines a next destination of a message based on the existence of one or more attributes of the message. In another example, a procedure can be defined that effects a determination based on the comparison of one or more attributes of the message to a reference value. In yet another example, a procedure can be defined that effects a determination based on pattern matching (e.g., regular expression matching). As would be appreciated, routing scripts can embody any of a variety of criteria-based procedures.

[1066] Routing scripts can specify a sequence of services that operate on either inbound or outbound messages for a service. As noted, in-transit services may themselves have routing

scripts requiring processing by other services. Therefore, the route calculation can be recursively defined based upon routing scripts specified by all services that interact with the message.

[1067] In one example, the sending service 310 may specify a routing script that requires a cross-reference mapping service to be included in the calculated route whenever sending service 310 sends a message to recipient service 360. In another example, recipient service 360 may specify a routing script that requires that any incoming request messages must first pass through a filter service to block messages from a list of sending services 310.

[1068] Routing scripts enable sending services to include services into the message route without having to explicitly specify the services in the message itself. Also, routing scripts enable recipient services 360 to require services 350 to be in the calculated route, regardless of the sending service's route specification.

[1069] In one embodiment, a routing script is embodied as a routing rule. A routing rule includes two parts: a condition and one or more resultant actions. The conditional part of a rule can be based on any elements or element attributes in a message's header. Additionally, content-based routing can be supported through conditional rules based on attributes of an element in a message's body and/or attachments.

[1070] Every rule should have at least one condition. Conditions include an operator and zero or more operands. Example operators include equals, notEquals, equalsOneOf, lessThan, greaterThan, and exists operators. In one embodiment, operators act on XML elements, XML attributes, or on other conditions.

[1071] From the standpoint of the element operators, XML elements contain either child elements or character data. Therefore, the operands for an element comparison both represent

the same type of content: either elements or character data. Character data can be in the form of a string, number, or date. Conditions involving elements that do not appear in the message will evaluate to false.

[1072] Attributes always have a type of character data, which can be string, number, or date. Many attributes are implicitly included in an XML document with default values. Therefore, an attribute identified in a condition can refer to either an explicit or implicit attribute. Conditions involving optional attributes that do not appear in the message will evaluate to false.

[1073] The usual boolean operators can combine conditions into more complex conditions. Condition operators act on other conditions. Example condition operators include AND, OR, XOR, and NOT condition operators.

[1074] The result of satisfying a rule's conditions is that an action will be triggered to modify the route for a message. Probably the most common result of a rule is to add one or more services into the route for a message. Several rule actions can be defined, including but not limited to an AddServiceAfter action, an AddServiceBefore action, an AddService action, a Redirect action, a ChangeTopic action, and a StopRuleEvaluation action.

[1075] In an AddServiceAfter action, a service (other than a final recipient service) can add a service after itself in the route. If this action appears more than once in the action list for a rule, the service is added such that the resultant service order is the same as the order of actions.

[1076] In an AddServiceBefore action, a service (other than a sending service) can add a service prior to itself in the route. If this action appears more than once in the action list for a rule, the services are added such that the resultant service order is the same as the order of actions.

[1077] The AddService action is identical to either the AddServiceAfter action or the AddServiceBefore action depending on the role the service has with respect to the message. For message senders, the services are added after the sender; for message recipients, the services are added prior to the recipient; for in-transit services, the services are added prior to the including service. This action generally enables rules that are useable by a service independent of role.

[1078] In a Redirect action, a service may wish to have the message redirected to another service in its place as the receiving service for the message. For example, a service may be a virtual service, and needs to redirect messages for that service to another service.

[1079] In a ChangeTopic action, if a first service includes a second service into the route or performs a redirect for a third service, then the first service can also change the topic of messages sent to the second service or the third service. The topic change will only apply to messages as delivered to the second service or the third service, and will revert to the original topic for subsequent services in the message route.

[1080] The StopScriptEvaluation action terminates the evaluation of subsequent scripts for the same service. Script evaluation will then continue for the next service in the route.

[1081] A service should maintain an evaluation sequence for the scripts associated with each role that the service can have with respect to a message. That sequence determines the order in which the scripts for that service are applied.

[1082] In one embodiment, scripts are evaluated in the following order: (1) scripts for the sender of the message; (2) scripts for services included by the scripts for the sender (this is recursive); (3) scripts for the recipients of the message, in the order of recipients in the message header; and (4) scripts for services included by scripts for the recipients (this is recursive).

[1083] When multiple scripts for a service include services into a route, the order of services in the route will follow the order of the scripts. That is, if script 1 inserts service A, and script 2 inserts service B, and if script 1 is evaluated before script 2, then service B follows service A in the route.

[1084] In one embodiment, routing scripts are evaluated only once during the initial calculation of the route for a message. The message header contains the basic information to initially construct a route, such as sending service 310 and recipient services 360. The message can also contain an explicit specification of a set of services to include in the route. Once the route is constructed from the header information, routing scripts are applied to further elaborate the route.

[1085] In an alternative embodiment, at least part of the message route is calculated after the physical routing of the message has begun. Dynamic routing is described in greater detail below in the context of physical and logical routing.

[1086] At the transport level, message interchange network 150 routes a message to a service by delivering the message through the Internet to a physical machine on which the service resides. That service operates on the message and, if the message is a request, returns a response message back through the Internet to message interchange network 150. The sequence of message deliveries and responses between message interchange network 150 and services represents the physical routing of a message.

[1087] Message interchange network 150 also provides a mechanism for a service to act on a message without the message being physically delivered to the service over the Internet. This mechanism is enabled through the logical routing of the message to the service. With logical

routing, a service can modify the routing of the message or modify the context of the message for delivery to the next service. Significantly, a service can be logically included in a message routing, without being included as part of the physical routing of the message.

[1088] In one embodiment, logical routing of messages is implemented through the specification of routing scripts. As described above, a service can define one or more routing scripts. These defined routing scripts are stored within message interchange network 150 and are processed to determine what routing behavior should occur when a message is logically routed to the service.

[1089] Logical routing can take place statically or dynamically. With static logical routing, a message is logically routed to all services prior to any physical routing. In other words, message interchange network 150 logically routes the message to all services prior to the physical delivery of a message to any services. This logical routing is represented by the sequential evaluation of the routing scripts that are defined by those services. As noted above, in one embodiment, the routing scripts are evaluated in the following order: (1) scripts for the sender, (2) scripts for the services included by the sender (recursive), (3) scripts for the recipient, and (4) scripts for the services included by the recipient (recursive).

[1090] In dynamic logical routing, the logical routing is not completed prior to the start of the physical routing. Rather, the logical routing takes place in sequence with the physical routing of the message. The relation between logical routing and physical routing is described in greater detail below.

[1091] As noted, message interchange network 150 delivers a message logically to every service participating in a message's routing. Of those services, some subset will also accept physical delivery of the message.

[1092] To illustrate this concept, consider an example where service A includes service B into the message route. Service A can include service B into the route either prior to itself in the route (provided service A is not the originator of the message) or after itself in the route. In either case, message interchange network 150 would logically route the message first to service A, which includes service B into the route. Message interchange network 150 then logically routes the message to service B, and after service B produces a response, message interchange network 150 logically returns the response to service A. The point at which service A physically receives the message depends on whether service A included service B prior or after itself in the route. If service A includes service B prior to itself in the route, then the order of physical delivery is first to service B then to service A. Conversely, if service A includes service B after itself into the route, then the order of physical delivery is first to service A then to service B. In the latter case, the response from B is not necessarily physically delivered back to service A. Rather, it may be only logically delivered back to service A.

[1093] Services to which a message is logically routed do not necessarily have to also physically receive the message. In the above example, service A could have been logically routed, with physical delivery only to service B. Consider the following scenario. Suppose service X includes service A into the route and service A includes service B into the route. The logical routing of the message would proceed from service X to service A to service B back to

service A back to service X. Service A can choose not to be included into the route for physical delivery, in which case the physical routing of the message is from service X to service B.

[1094] In general, the act of routing a message (physically or logically) to a service can be thought of as an invocation of the service. When a service includes another service into the route of a message, the including service is effectively invoking the included service. The invocation of a service does not necessarily imply the physical delivery of information to the invoked service. The logical routing of a message is then the logical invocation of services. A route that includes a progression of services including other services can effectively be modeled as a progression of invocations.

[1095] In logical routing, each service is not only able to manage the inclusion of other services into the route but is also able to manage the context of those inclusions. From the standpoint of invocations, an invoking service is able to set the context for the invocation. An invoked service can also set the context of its return.

[1096] It is a feature of the present invention that message interchange network 150 effects context management on behalf of invoking services. As noted, while an invoking service can be logically included in a message routing, it need not be included as part of the physical routing of the message. In general, message interchange network 150 persistently stores contexts of a message, thereby enabling proper restoration of contexts upon return from an invocation.

[1097] In various embodiments, invocation context can include such information as the identity of the invoker service, arguments to the invoked service, a session identifier for the message, a topic for the message, billing responsibility for the invocation, or any other information that can be used by the invoked service. When a service is invoked, it receives a

context from its invoker. That invoked service can then modify, if desired, parts of the context when invoking yet another service. Upon return from an invocation, message interchange network 150 automatically restores the context of an invoker to the state prior to the invocation. In one embodiment, the context of an invocation is included in the header information of the message delivered to a service. In another embodiment, the context of an invocation is based on one or more parts of a message.

[1098] Responses from invocations also contain context, such as the completion or error status of the response. When a service receives a returned response from an invocation, the context of the invoker is augmented by the context of the response. That service can then optionally modify the response context when returning to its invoker. It should be noted that just as logical routing can occur either statically or dynamically, context propagation (for both invocation and response) can also occur either statically or dynamically.

[1099] The ability of services to manage the context of their invocations provides a way for services to assume a large degree of control over the scope of their participation in handling a routed message. For example, a service can choose to isolate its invoked services from being aware that the invocation occurs as part of a broader message routing. A service can choose to assume responsibility for charges incurred by subsequent nested invocations. A service can also choose whether or not to expose a received error condition up the invocation chain.

[1100] Message interchange network 150 can also use the invocation contexts to track messages and determine responsibility for invocations. Based on invocation contexts, message interchange network 150 can reconstruct the history of a message necessary for auditing and billing, as well as for reporting.

[1101] Having described a framework for logical routing and invocation of services, the description of the physical routing process is continued with reference to FIG. 3. In this example, it is assumed that static logical routing has produced a route for the message.

[1102] After deriving the route for a message, message interchange network 150 validates the route. There are numerous conditions that can cause a route to be invalid. For example, there may be routing permission violations or a service may be currently disallowed by message interchange network 150 due to, for example, the non-payment of usage bills.

[1103] If the route is determined to be invalid, then message interchange network 150 rejects the posted message and may return an error to the sending service 310 either in the response to the posting call or in an error message.

[1104] Message interchange network 150 routes the message to all services in the calculated route. In the event of failure at any stage in the routing, message interchange network 150 aborts the message routing and, if the original message was a request, returns an error message back to sending service 310. For messages with multiple recipients, an error in the routing to one recipient will not necessarily affect routing to other recipients. Errors during routing can occur due to several circumstances. For example, a message may fail to reach a recipient within the expiration time for the message, a service may fail to return a reply to a delivered message, or a service may return an error status for a message.

[1105] Message interchange network 150 sequentially delivers a message to each service identified in the message route. In FIG. 3, this process is illustrated as a flow of the posted message through message routing element 340, and on to one or more services 350. In one embodiment, message interchange network 150 includes all of the message documents in the

message delivered to service 350, even if service 350 only expects to operate on one document or documents of a particular content-type. Service 350 would ignore documents that it does not expect.

[1106] As noted, message interchange network 150 invokes in-transit services in the same way as delivering a message to any sending or recipient service. In general, a service does not necessarily need to be aware whether it is being invoked as an in-transit service or as a recipient service.

[1107] After processing the message, service 350 sends the results in a response message back to message interchange network 150. If service 350 is unable to produce a valid result for its operation on a message, then service 350 may return an error code in its response message to message interchange network 150. If service 350 fails to respond to a received message, the message will ultimately expire, and message interchange network 150 may return an error message back to the sending service 310.

[1108] Upon receipt of the response message from service 350, message interchange network 150 then routes the message to the next destination (e.g., another service 350) on the routing list. After passing through each of the intermediate destinations on the routing list, the message is then stored in queue 334 for recipient service 360. It should be noted that queues can also be associated with in-transit services 350. For simplicity, these queues are not shown in the illustrated embodiment of FIG. 3.

[1109] Application 372 in recipient service 360 can retrieve the message from queue 334 via message poll interface 328. In the poll mode, application 372 periodically issues a call to message poll interface 328 to ask for any waiting messages. If there are queued messages

waiting for delivery to that service, then the one or more messages are returned in a reply. When making poll requests, a service can provide selectors on the messages to fetch. For example, a service can retrieve messages based upon the sender, message type, topic, etc.

[1110] In an alternative embodiment, message delivery is enabled through a push mode. In the push mode, the recipient would have its own server to which message interchange network 150 can send messages. A service can specify a maximum number of tries, as well as the retry interval, for message interchange network 150 to send the message to the service before aborting pushed delivery. A service to which message interchange network 150 pushes messages can also optionally poll for messages.

[1111] In the push mode, a service can also specify a delivery window in which it will accept pushed messages. For example, a service might only accept messages between the hours of 1 AM and 2 AM.

[1112] As further illustrated in FIG. 3, a response message can be posted by recipient service 360 to message post interface 326. The message is then routed through one or more services 350 prior to being stored in message queue 332. The message can then be retrieved by sending service 310 through message poll interface 322. As noted, the return path would not necessarily match the forward path.

[1113] In one embodiment, the sender of a message can specify a time by which the routing of a message must complete. If that expiration time passes before delivery of the message to its final destination, message interchange network 150 will abort further routing of the message. In one embodiment, the final destination for a request message is the delivery of a response message back to the original request sender. In this embodiment, senders of response messages

cannot specify an expiration since the response is considered part of the routing of the original request message. If a request message expires, message interchange network 150 will return an error response back to the sender. If the sender does not specify an expiration time, then a default message expiration time (e.g., 48 hours) can be used. It should be noted that message expiration is not the same as document expiration. Document expiration is associated with a specific document and indicates how long the document's information is valid.

[1114] As part of the message delivery process, message interchange network 150 logs all posted messages, including invalid messages. For each message, message interchange network 150 logs relevant information to track the history of the message. Message interchange network 150 also maintains a correlation between messages. That is, for request messages, message interchange network 150 associates the log of the response message(s) with the log of the request message.

[1115] In one embodiment, logged information can include the message header, the calculated route for the message (unless route calculation fails), the status of route validation, the size of the message at each stage of the route, and the routing history, including the status for each service along the message's route. The status values for each service depends on the role of the service.

[1116] Message interchange network 150 correlates all messages for the same message transaction. That is, for request messages, message interchange network 150 associates the log of the response message(s) with the log of the request message. Similarly, if a message causes an error message, then message interchange network 150 associates the log of the generated error message with the log of the original message.

[1117] Having described a general framework of operation of message interchange network 150, an example message sequence is provided to illustrate the concepts described above. In the example message sequence illustrated in FIG. 4, a message is routed from sender 402 to recipient 404. During this message routing, the message is also passed through in-transit services 406, 408, 410, and 412. The specific operation of services 406, 408, 410, and 412 will become apparent through the description of the message operation sequence.

[1118] In the illustrated example, sending service 402 represents a business named MyBiz.com. MyBiz.com desires to send a purchase order to The Acme Company, which operates as recipient service 404. Before being delivered to The Acme Company, the purchase order message is sequentially routed by message interchange network 150 through Transmatics (XSLT) service 406, XpandiCo service 408, Transmatics (replaceElement) service 410, and KeepEmOut.com service 412.

[1119] In general, services 406, 408, 410, 412 are operative to perform transformations, enrichment, and filtering functions on the purchase order message. In the present example, Transmatics (XSLT) service 406 is operative to transform address data in the message body, Xpandico service 408 is operative to augment a zip-code address, Transmatics (replaceElements) service 410 is operative to relate sets of values from disparate data sources, and KeepEmOut.com service 412 is operative to filter messages for The Acme Company.

[1120] As noted, services 406, 408, 410, 412 can be added to a routing list based upon routing scripts that are created for one or more of sending service 402, recipient service 404, or in-transit services 406, 408, 410, 412. In the present example, services 406, 408, and 410 are added to the route list in accordance with routing scripts defined for sending service 402, while

service 412 is added to the route list in accordance with routing scripts defined for recipient service 404. An example embodiment of the message sequence between services in the routing list is now provided.

[1121] As noted, in one embodiment, a message includes a header element, a body element and/or attachments. In general, the header element includes the basic delivery information for the message, and the body element/attachments includes the document(s) the sender is sending to the recipient. As would be appreciated, further message elements can be defined. For example, in another embodiment, a message would include header, body, routing, and route trace elements. Here, the routing element includes a listing of a sequence of services that should operate on the message prior to delivery to the recipient, and the route trace element includes the routing history information.

[1122] Appendix B.1 illustrates an example of a message that is posted by MyBiz.com to message interchange network 150. As illustrated, the message includes a header element at lines 4-17 and a body element at lines 18-28.

[1123] The header element at lines 4-17 includes elements To, From, and Expiration. To element also includes two Service elements, one of which includes the element Arguments. As illustrated, one of the Service elements identifies the recipient service "TheAcmeCompany/Supply," the service that is representative of recipient 404 in FIG. 4. The element Arguments provides parameters that further specify what the "TheAcmeCompany/Supply" service is to perform (i.e., ProcessPurchaseOrder).

[1124] As illustrated, at line 14, the From element includes a service identifier for MyBiz.com (i.e., “386b4520f489c217”). Finally, at line 16, the Expiration element specifies a time by which the message should complete its routing.

[1125] The body element of the message is located at lines 18-28. In this message element, the purchase order data is included at lines 21-25. As will be described below, the purchase order data can be transformed and enriched by services in the routing list.

[1126] As noted, the sender can specify an explicit sequence of services that should operate on the message. The sender can also implicitly include services in the routing list for a message through the specification of routing scripts associated with the defining service. In the example message of Appendix B.1, no services are explicitly identified in the message. Rather, the services are included based on routing scripts that are defined by the various services. In the present example, the routing list includes, in order, Transmatics (XSLT) service 406, XpandiCo service 408, Transmatics (replaceElement) service 410, and KeepEmOut.com service 412.

[1127] The first service on the routing list is the Transmatics (XSLT) service 406. In the present example, service 406 is included in the routing list based upon a routing script defined for sender 402. Transmatics (XSLT) service 406 is a transformation service that uses the Extensible Stylesheet Language Transformations (XSLT) language.

[1128] Generally, XSLT is a language for transforming XML documents into other XML documents. XSLT is designed for use as part of the XSL, which is a stylesheet language for XML. In addition to XSLT, XSL includes an XML vocabulary for specifying formatting. XSL specifies the styling of an XML document by using XSLT to describe how the document is transformed into another XML document that uses the formatting vocabulary.

[1129] In the present example, Transmatics (XSLT) service 406 is operative to transform data in the purchase order message. This transformation is illustrated in Appendix B.2, which shows the purchase order message as it is delivered to XpandiCo service 408 after being processed by Transmatics (XSLT) service 406. More specifically, the transformation operation is illustrated upon comparison of the Address element at line 25 of the message in Appendix B.1 to the Address element at lines 24-28 of the message in Appendix B.2. In this transformation process, an Address element including the entire address has been transformed into an Address element that includes further child elements (i.e., Street, City, and Zip) directed to components of the Address.

[1130] As can be appreciated, this transformation process can be used to ensure that the purchase order to be delivered to a particular recipient service 404 is placed in the proper format. If other recipients require a purchase order to appear in different formats, then further transformation operations can be defined and implemented by a service attached to message interchange network 150.

[1131] After the message is processed by Transmatics (XSLT) service 406, it is then sent to the next service on the routing list. In this example, the next service on the routing list is XpandiCo service 408.

[1132] As illustrated in Appendix B.2, the message delivered to XpandiCo service 408 includes Header and Body elements that are similar to the corresponding elements in the message that was posted by MyBiz.com.

[1133] The Header element in the message delivered to XpandiCo service 408 also includes a Session element and a Token element. The Session element at line 4 includes a unique session

identifier (i.e., “34b9f6d0-89eb-b5e1-0022-a376bf41c165”) that is assigned by message interchange network 150 to the message. This unique session identifier enables tracking of the message through the routing network. The Token element at line 5, on the other hand, is a unique reference identifier (i.e., “84e309b38c56a18cb9835203”) that enables message interchange network 150 to uniquely identify a delivered message in the routing network. This unique identifier serves as a message reference in response messages.

[1134] As further illustrated in Appendix B.2, the To element, at lines 6-12, includes the service identifier (i.e., “3340f32c035d7499”) for XpandiCo as well as the enrichment operation that is desired to be invoked. Specifically, at line 7, the purchase order message identifies the “zipPlus4” operation. The “zipPlus4” operation is an example of an enrichment operation, wherein a 5-digit zip code is expanded to a 9-digit zip code.

[1135] After the “zipPlus4” enrichment operation is completed, XpandiCo service 408 returns a response message to message interchange network 150. The returned response message is illustrated in Appendix B.3. As appears at line 19 of the message, the Zip child element of the Address element has been modified to include a 9-digit zip code.

[1136] After XpandiCo service 408 returns a response message to message interchange network 150, the message is then forwarded by message interchange network 150 to the Transmatics (replaceElement) service 410. In general, Transmatics (replaceElement) service 410 provides a cross-reference ID mapping function. This function provides the ability to relate sets of values from disparate data sources to each other. These related values are stored persistently by the cross-reference service 410 and can be substituted for each other whenever a message travels from one of the sources to another.

[1137] This cross-reference function is especially useful for relating key or ID values of two independent data sources. For example, consider a scenario where two different services independently maintain a customer contact list for the same customer. In this case, it would be difficult for each service to process a change message from the other without a cross-reference map because the receiving service would have no explicit way of knowing which record to update. By using cross-reference service 410, however, the receiving service's record ID could be automatically substituted for the sender's record ID, thereby enabling the receiving service to easily process the message and update the correct record.

[1138] In one embodiment, a cross-reference map actually stores a set of records for each service sharing a cross-reference map. Each record can include an arbitrary number of name/value pairs, although one of them is identified as the key. Records from different services can then be related to each other in one of two ways. In one method, different services can be related to each other explicitly by relating two key name/value pairs from different services. In another method, services can be related to each other automatically by having the cross-reference service 410 process the response to a message as well as the originating message.

[1139] As would be appreciated, the cross-reference map can be queried later by specifying a specific service, key name, and key value. The extracted name/value pairs from all related records can then be used to transform a message.

[1140] In the example of FIG. 4, Transmatics service 410 is being called upon to perform a replaceElement operation on the CustomerNumber. After the CustomerNumber is replaced with a number that is recognized by The Acme Company, Transmatics service 410 returns a response message to message interchange network 150. As illustrated in Appendix B.5, the customerID

element at line 23 has been changed to “BX0045012” from the previous customerID value of “3489.”

[1141] After the replaceElement operation has been performed by Transmatics service 410, the message is then forwarded to KeepEmOut.com service 412. KeepEmOut.com service 412 performs a message filtering service by ensuring that The Acme Company will only receive messages from authorized entities. As such, KeepEmOut.com service 412 can be included into the routing list based on a routing script that has been defined by The Acme Company.

[1142] If the message passes through KeepEmOut.com service 412, then the message is finally delivered to The Acme Company service 404. An example message as delivered to The Acme Company is illustrated in Appendix B.4. An example of a response message by The Acme Company to MyBiz.com is also illustrated in Appendix B.5.

[1143] As thus described, message interchange network 150 enables flexible interaction with third-party services. The connection of these third-party services to an open platform enables enterprise customers to flexibly define their message transaction framework. As will be described in greater detail below, message interchange network 150 can also be used to enable flexible interaction between enterprise customers and their ASPs.

[1144] In general, the distribution of an enterprise’s information among ASPs requires a security framework by which each ASP can authenticate access to that enterprise’s information by another ASP. In a point-to-point-messaging model, the conventional approach is to embed the authentication information (encrypted) in the message itself or in some challenge protocol. In a collaborative environment, this authentication approach suffers several disadvantages. First, the originator of a message may be an ASP acting on behalf of an enterprise, and will therefore

be unaware of the authentication credentials to present to the recipient ASP. Second, the authentication credentials are applicable only to the recipient ASP and do not authenticate access of other services in the message's route. Third, many systems use standard HTTPS-based authentication, which requires synchronous invocation of the recipient. Polling for messages would therefore be disallowed. Finally, services along a message's route that operate on the message's content should not have visibility to the authentication information.

[1145] In accordance with the present invention, message interchange network 150 authenticates each service that participates in a message's route. Message interchange network 150 can then provide an authentication token that substitutes for an enterprise's authentication credentials. The proposed token is an identifier (GCID) that represents an authenticated service to the enterprise's account with a specific ASP. An ASP that receives a GCID in a message only needs to authenticate message interchange network 150 since message interchange network 150 has already authenticated the message sender. The ASP can either have an internal map between the GCID and enterprise account identifier, or can require the account identifier to accompany the GCID.

[1146] If the ASP does not want to maintain its own internal mapping between an enterprise's GCID and the enterprise's account, then the ASP would provision the enterprise's account identifier to message interchange network 150. Message interchange network 150 would then insert the account identifier with the GCID in any subsequent messages from the enterprise to ASP 620.

[1147] For both an enterprise and ASP to trust a GCID provided by message interchange network 150, the GCID should be validated by both organizations as representing the

enterprise's account on the ASP. Setting up this trust relationship is accomplished through a provisioning process that is completed prior to the routing of messages between the enterprise and the ASP. The process of provisioning a trusted GCID in message interchange network 150 involves two authentication steps: the enterprise first authenticates itself to message interchange network 150 and receives a GCID, then the enterprise authenticates itself to the ASP and associates the GCID with the enterprise's account.

[1148] FIG. 6 illustrates an embodiment of the provisioning process. It should be noted that in the illustrated embodiment of FIG. 6, it is assumed that the enterprise and the ASP have already registered with message interchange network 150.

[1149] The provisioning process begins at step (1) where ASP 620 registers with message interchange network 150, alerting message interchange network 150 that its service requires provisioning. In this process, ASP 620 provides a URL to message interchange network 150. The provided URL enables an enterprise user 610 to logon to the ASP's website for provisioning.

[1150] At step (2), enterprise user 610 logs in to the message interchange network's web site. This authenticates enterprise user 610 with message interchange network 150. Enterprise user 610 then registers for a GCID at step (3). This GCID will represent an authenticated entity that can both send and receive messages through message interchange network 150. At step (4), enterprise user 610 selects a service from the message interchange network directory representing an ASP with whom the enterprise has an account.

[1151] At step (5), message interchange network 150 then automatically redirects enterprise user 610 to a website at ASP 620. In this process, message interchange network 150 sends the enterprise's GCID and a provisioning token to the web page.

[1152] In one embodiment, message interchange network 150 passes a GCID, a provisioning token, and a return URL as a query string in the ASP's URL using three parameters: serviceID - an 8 byte hex number representing the GCID registered by an enterprise user; token - an 8 byte hex number representing a unique token to be returned in the confirmation message from the ASP back to message interchange network 150; and returnUrl - a URL string representing where the user should be redirected to after completing the authentication and provisioning process at the ASP website.

[1153] An example of a redirect URL would be:

`https://www.asp.com/gc?serviceID=345023ba2300d353&token=459012c9a78b90af
&returnURL=https://www.grandcentral.com/provisioning/auth.jhtml`

[1154] On the ASP's website, enterprise user 610 logs in to ASP 620 at step (6), thereby being authenticated by ASP 620. ASP 620 will then recognize the validity of the GCID as representing enterprise 610. At step (7), the ASP's website can optionally allow enterprise user 610 to select events (if any) for which the ASP's service should send notification messages to the enterprise's GCID.

[1155] After authenticating enterprise user 610 and allowing enterprise user 610 to select events, the ASP website redirects enterprise user 610 back to the message interchange network website at step (8). Next, at step (9), ASP 620 posts a message to message interchange network

150 indicating that ASP 620 has accepted the enterprise user's authentication of the GCID. In one embodiment, the message contains both the GCID and the provisioning token originally sent to the ASP's webpage by message interchange network 150. In general, this back-channel confirmation, as opposed to returning a confirmation in a URL redirect back to message interchange network 150, prevents spoofing of the confirmation.

[1156] In one embodiment, ASP 620 posts the confirmation message to message interchange network 150 using the message interchange network messaging protocol. Message interchange network 150 would expect specific values in some of the message elements. An example of a confirmation message (showing only the relevant message elements) is shown below. As would be appreciated, the actual values of the arguments and document content would be dependent on the specific provisioning circumstances.

```
<Message>type="request">
  <Header>
    ...
    <To>
      <Service>name="grandcentral/provisioning">
        <Arguments>
          <![CDATA[
            <token>459012c9a78b90af</token>
            <gcid>63801c78327fd901</gcid>
            <operation>name="setPermission">
              <parameter>include</parameter>
            </operation>
            <operation>name="setCookie"/>
          ]]>
        </Arguments>
      </Service>
    </To>
    ...
  </Header>
  <Body>
    <Document>id="COOKIE">content-type="text/xml">encoding="none">
      <![CDATA[
        <!-->insert>enterprise>account>information>-->
      ]]>
    </Document>
```

</Body>
</Message>

[1157] The values for the GCID and token parameters in the arguments should be the values received via the web interface redirection. The cookie parameter (and the document in the message body) can be omitted from the confirmation message if ASP 620 maintains the map between the GCID and enterprise account identifier. Otherwise, the contents of the document element should be whatever text or XML ASP 620 requires to determine the enterprise's account identifier.

[1158] In general, a service can register a block of text or XML (referred to as a cookie) that message interchange network 150 is to include with a message whenever a message is delivered to the service. This cookie is analogous to cookies in web browsers. An example of a cookie is account information that a service requires when receiving a message from particular sources. This eliminates the need for the sender of a message to include, or even be aware of, this information when sending the message.

[1159] If enterprise user 610 fails authentication by ASP 620, ASP 620 should not send a confirmation message to message interchange network 150. ASP 620 should, however, provide immediate feedback to enterprise user 610 on the web page that authentication failed.

[1160] After processing the confirmation message, message interchange network 150 returns a status response to ASP 620. The status element in the response message will be absent in the case of a success and otherwise will contain an error code.

[1161] In one embodiment, message interchange network 150 must receive the confirmation message from ASP 620 within 24 hours after enterprise user 610 initiates provisioning a GCID.

If message interchange network 150 does not receive the confirmation within that time period, then the provisioning token will expire. If message interchange network 150 receives a confirmation message after the token has expired, message interchange network 150 would not recognize the confirmation and would return an error back to ASP 620. If no confirmation is received, message interchange network 150 would not notify either ASP 620 or enterprise user 610 about the expiration of the provisioning token.

[1162] In an alternative embodiment of the provisioning process, ASP 620 provides a launch point from its own website for provisioning. ASP 620 could then redirect a user originally at the ASP site over to the message interchange network site to register for a GCID. In this embodiment, step (8) would not be necessary after message interchange network 150 redirects the user back to the ASP's site.

[1163] It is a feature of the present invention that the provisioning process can set up a trust relationship between the enterprise and the ASP. When ASP 620 returns the provisioning confirmation message to message interchange network 150, ASP 620 is indicating that its service trusts the enterprise's GCID for sending messages and trusts all of the enterprise's GCIDs for receiving messages. Similarly, by this provisioning process, the enterprise is indicating trust of ASP 620 for sending and receiving messages.

[1164] It should be noted that if an ASP does not wish to expose the notion of accounts (i.e., messages to the ASP are not account-specific), or if the ASP does not wish to exploit the message interchange network's authentication framework, then the ASP may choose not to require authentication of an enterprise's GCID to a specific account in the ASP, and this provisioning process is not required.

[1165] If ASP 620 wants to revoke the authenticated association between an enterprise's GCID and the enterprise's account in ASP 620, ASP 620 can simply cease to recognize the association between the GCID and account. In this case ASP 620 would return an error response whenever it receives a message from that GCID. ASP 620 can also use the message interchange network website to unset permissions granted by ASP 620 to the enterprise. Message interchange network 150 would then not allow the enterprise's GCID to send messages to ASP 620.

[1166] Alternatively, ASP 620 can send a request message to message interchange network 150 to revoke the previously provisioned permissions and cookie. In one embodiment, this message would contain the following information:

```
<Message>type="request">
  <Header>
    ...
    <To>
      <Service>name="grandcentral/provisioning">
        <Arguments>
          <![CDATA[
            <gcid>63801c78327fd901</gcid>
            <operation>name="clearPermission">
              <parameter>include</parameter>
            </operation>
            <operation>name="clearCookie"/>
          ]]>
        </Arguments>
      </Service>
    </To>
    ...
  </Header>
  <Body/>
</Message>
```

[1167] The effect of this revocation message is that message interchange network 150 automatically revokes all permissions granted by ASP 620 to the enterprise, and message

interchange network 150 removes the mapping of the enterprise's GCID to ASP's cookie in the message interchange network registry.

[1168] After processing the revocation message, message interchange network 150 returns a status response to ASP 620. The status element in the response message is absent in the case of a success and otherwise will contain an error code.

[1169] After an enterprise provisions a trusted GCID between an enterprise and an ASP's service, the GCID becomes a mapped service representing the ASP's service specific to the enterprise's account. The following scenarios illustrate how this GCID would be used in various messaging situations.

[1170] First, assume that the enterprise wishes to send a message to the ASP specific to the enterprise's account with the ASP. The enterprise should then send a message to its mapped service for the ASP:

```
<Message>
  <Header>
    <From>
      <Service>gcid="enterprise"/>
    </From>
    <To>
      <service>gcid="enterprise/mappedService"/>
    </To>
  </Header>
</Message>
```

[1171] This would be translated by message interchange network 150 for delivery to the ASP:

```
</Message>
```

```
<Header>
  <From>
    <service>gcid="enterprise/mappedService"/>
  </From>
  <To>
    <Service>gcid="asp/service">
      <Cookie><![CDATA[account=12345]]></Cookie>
    </Service>
  </To>
</Header>
</Message>
```

[1172] It should be noted that the enterprise could send the message directly to the ASP rather than to the mapped service, and the cookie would still be inserted correctly by message interchange network 150. The advantage for the enterprise in sending the message to the mapped service rather than directly to the ASP is to facilitate a consistent tracking of invocations of the ASP whether from the enterprise or from another ASP on behalf of the enterprise.

[1173] In a second scenario, assume that an enterprise has accounts with more than one ASP. When integrating the exchange of information between these ASPs, the enterprise may allow one ASP to send requests or notifications to another ASP on behalf of the enterprise. With mapped services, the ASPs do not actually need to know each other specifically. A first ASP would send a message to the enterprise's service mapped to the second ASP:

```
<Message>
  <Header>
    <From>
      <Service>gcid="asp1/service"/>
    </From>
    <To>
      <service>gcid="enterprise/mappedService"/>
    </To>
  </Header>
</Message>
```

[1174] This would be translated by message interchange network 150 for delivery to the second ASP:

```
</Message>
  <Header>
    <From>
      <service>gcid="enterprise/mappedService"/>
    </From>
    <To>
      <Service>gcid="asp2/service">
        <Cookie><![CDATA[account=12345]]></Cookie>
      </Service>
    </To>
  </Header>
</Message>
```

[1175] Note that the message as delivered to the second ASP hides the identity of the first ASP. In effect, message interchange network 150 effects a virtual proxy service in mapping between services. Accordingly, if the message is a request, then the second ASP's response message would be returned to the enterprise's mapped service and redirected to the first ASP.

[1176] Having described the flexibility in the operation of message interchange network 150, an embodiment of message interchange network 150 is now provided with reference to FIG. 5. FIG. 5 illustrates a message routing network 500 that includes clients 510, points of presence (POP) 520-n, and data center 530.

[1177] Clients 510 are generally assigned to a particular POP 520-n. POPs 520-n include a plurality of servers 522, a message queue 524, and a directory cache 526. Directory cache 526 includes profile information on clients 510. The content of directory cache 526 is based on profile database 532 in data center 530. Message queue 524 is operative to store messages that are posted by clients 510 and to store messages that are awaiting delivery to clients 510. Finally,

servers 522 are lightweight servers that process HTTP transactions. Servers 522 are operative to receive messages posted by clients 510, and to send messages to clients 510. As noted, clients 510 can have messages pushed to them or can retrieve messages as part of a polling process.

[1178] Routing within message routing network 500 is generally enabled through routers 531 in data center 530. In one embodiment, routers 531 are operative to pull messages from message queues 524 in POPs 520-n, perform route calculations, and post messages to message queues 524 that are located in a POP 520-n that is assigned to the destination client 510.

[1179] Headers of messages that are routed by routers 510 are stored in message database 533. Message database 533 retains updated states of message objects that are routed through message routing network 500.

[1180] Also included within data center 530 is profile database 532, data warehouse 535, and web interface component 534. Profile database 532 stores profile information for registered clients 510. The profile information includes routing rules that have been defined for the various clients 510. Data warehouse 535 stores a collection of log information relating to messages that are processed by message routing network 500. Finally, user interface component 534 enables clients 510 to enter profile information into profile database 532, access reports in data warehouse 535, and perform directory lookups to find other connected services.

[1181] While the invention has been described in detail and with reference to specific embodiments thereof, it will be apparent to one skilled in the art that various changes and modifications can be made therein without departing from the spirit and scope thereof. Thus, it is intended that the present invention cover the modifications and variations of this invention provided they come within the scope of the appended claims and their equivalents.